

I. REAL PARTY IN INTEREST

The subject application is owned by Sun Microsystems, Inc., a corporation organized and existing under and by virtue of the laws of the State of Delaware, and having its principal place of business at 4150 Network Circle, Santa Clara, CA 95054, as evidenced by the assignment recorded at Reel 015184, Frame 0027.

II. RELATED APPEALS AND INTERFERENCES

No other appeals or interferences are known which would directly affect or be directly affected by or have a bearing on the Board's decision in this appeal.

III. STATUS OF CLAIMS

Claims 1-20 are pending in the present application. Claims 1-5, 7-14, and 16-20 stand finally rejected and are the subject of this appeal. A clean copy of claims 1-5, 7-14, and 16-20, as on appeal (incorporating all amendments), is included in the Appendix hereto.

Claims 6 and 15 were objected to as being dependent on a rejected base claim, but were otherwise deemed allowable if rewritten in independent form to include all of the limitations of the base claim and any intervening claims, and are not a subject of this appeal.

IV. STATUS OF AMENDMEMNTS

No amendment to the claims has been filed subsequent to the final rejection. The Appendix hereto reflects the current state of the rejected claims.

V. SUMMARY OF THE CLAIMED SUBJECT MATTER

Independent claim 1 is directed to a system comprising a plurality of nodes (e.g., Fig.'s 1 and 6A, 140A, 140B; paragraphs 0057-0062, specification). Each node includes an active device (e.g., Fig 6, 142AA; paragraphs 0059-0060, specification) and a memory subsystem (e.g., Fig.'s 5B, 6, 144AA; paragraphs 0054-0056 and 0060-0062). An active device in a node is configured to generate a global address and translation information identifying a translation function (e.g., Fig. 8, 212, 222; paragraph 0067), and wherein the global address identifies a coherency unit (e.g., paragraph 0033). The memory subsystem is configured to select the translation function in response to the translation information (e.g., paragraph 0010; Fig. 9; paragraph 0071) and to perform the translation function on the global address to generate a physical address of the coherency unit within the memory subsystem (e.g., paragraph 0010; Fig. 7; paragraph 0066). An additional memory subsystem (e.g., Fig. 6A, 144AA) in an additional node (e.g., Fig. 6A, 140B) is configured to store the translation information identifying the translation function used in the node (e.g., Fig. 6A, 222B; paragraphs 0057-0058), wherein in response to a request for access to the coherency unit, the additional memory subsystem is configured to send the translation information to the node (e.g., paragraphs 0041-0044 and 0058-0062; Fig. 6A and 6B).

Claim 10 is a method for use in a system comprising a plurality of nodes (e.g., Fig.'s 1 and 6A, 140A, 140B; paragraphs 0057-0062, specification), each node including an active device (e.g., Fig 6, 142AA; paragraphs 0059-0060, specification) and a memory subsystem (e.g., Fig.'s 5B, 6, 144AA; paragraphs 0054-0056 and 0060-0062). The method (e.g., Fig.'s 3, 7, and 9; paragraphs 0047, 0065-0066, and 0071-0073) comprises an active device in a node generating a global address and translation information identifying a translation function (e.g., Fig. 8, 212, 222; paragraph 0067), wherein the global address identifies a coherency unit (e.g., paragraph 0033). A memory subsystem included in the node selects the translation function in response to the translation information (e.g., paragraph 0010; Fig. 9; paragraph 0071) and performs the selected translation function on the global address to generate a physical address of the coherency

unit within the memory subsystem (e.g., paragraph 0010; Fig. 7; paragraph 0066). An additional memory subsystem (e.g., Fig. 6A, 144AA) included in an additional node (e.g., Fig. 6A, 140B) stores the translation information identifying the translation function used in the node (e.g., Fig. 6A, 222B; paragraphs 0057-0058). In response to a request for access to the coherency unit, the additional memory subsystem sends the translation information to the node (e.g., paragraphs 0041-0044 and 0058-0062; Fig. 6A and 6B).

VI. GROUND S OF REJECTION TO BE REVIEWED ON APEAL

1. Claims 1-5, 7-8, 10-14, and 16-20 are rejected under 35 U.S.C. § 103(a) as being unpatentable over Chi, U.S. Patent 5,940,870 in view of Ang, U.S. Patent 6,678,799.

2. Claim 9 is rejected under 35 U.S.C. § 103(a) as being unpatentable over Chi in view of Ang and in further view of Arimili, U.S. Patent Application Publication 2002/0112124.

VII. ARGUMENT

A. Claims 1-2, 4-5, and 7-8

The Examiner rejected claims 1-2, 5, and 7-8 as being unpatentable over Chi in view of Ang under 35 U.S.C. § 103(a). Appellants respectfully traverse this rejection in light of the following remarks.

Independent claim 1 recites:

“a plurality of nodes, wherein each node includes an active device and a memory subsystem coupled to the active device;

wherein an active device in a node of the plurality of nodes is configured to generate a global address and translation information identifying

a translation function, wherein the global address identifies a coherency unit;

wherein a memory subsystem included in the node is configured to select the translation function in response to the translation information and to perform the translation function on the global address to generate a physical address of the coherency unit within the memory subsystem;

wherein an additional memory subsystem included in an additional node of the plurality of nodes is configured to store the translation information identifying the translation function used in the node, wherein in response to a request for access to the coherency unit, the additional memory subsystem is configured to send the translation information to the node” (Emphasis added).

Appellant respectfully submits that the prior art references, taken singly or in combination, fail to teach or suggest “translation information identifying a translation function”, “[selecting] the translation function in response to the translation information”, “[performing] the selected translation function”, or “storing the translation information identifying the translation function” as recited in combination with the other features of claim 1.

In the final office action, the Examiner contends that Chi teaches an active device configured to generate a global address and translation information identifying a translation function and a node configured to perform the translation function in Fig. 8. The Examiner further contends that Chi teaches an additional node configured to store the translation information identifying the translation function in Fig. 8, #114 and #110, and col. 5, lines 27-46. Applicant respectfully disagrees with the Examiner’s characterization.

In col. 5, lines 27-46, Chi describes an address translation from the source node to the destination node:

“FIG. 8 shows the address translation from the source node to the destination node. The processor bus address 100 at the source node consists of two parts: an AMT Index 102 and an Offset 104. The AMT Index 102 is used to retrieve an entry from an Address Mapping Table (AMT) 106. Each table entry 108 contains a destination node ID for an access request on the processor bus. The number of entries in the AMT 106 is not necessarily equal to the number of the address partitions. For instance, in one embodiment, the AMT has 64 entries and the number of address partitions is 16. In that case, the AMT Index 102 is 6 bits while only the most significant 4 bits of these 6 bits indicate the partition number 110. In the AMT 106, the four consecutive entries corresponding to a given partition number are filled with the same destination node. Note that the node ID from the AMT 106 is the physical node ID which is unique in the whole system. The partition number 110 is the **logical node ID** within each application, such that there may be more than one application running on the system.” (emphasis added).

Chi further states, in column 5, lines 55-65:

“The global address 112 on the interconnect is composed of the node ID 114 and the processor bus address 110 at the source node. This global address is part of the header of the packet sent from the source node to the destination node. At the destination node, the node ID is stripped off and the partition number is replaced with 0, such that the translated address 116 falls in the local space of the destination node. Finally, the private memory size 118 at the destination node (or the starting address of the near global memory of the destination node) is further added to produce the

processor bus address 120 to be used at the destination node.” (Emphasis added).

As evidenced by the above two citations, item 110 is taught as either a partition number that is a logical node ID, or a processor bus address. Chi provides no teaching or suggestion that item 110 provides any information that would identify a translation function. As evidenced in the second of the two citations, item 114 is taught by Chi as a node ID. Chi provides no other teaching regarding node ID 114, much less any teaching that node ID 114 identifies a translation function. Furthermore, Chi provides no other teaching or suggestion of “translation information identifying a translation function”, nor does Ang provide any teaching or suggestion that, taken either singly or in combination with Chi that would remedy this deficiency. Finally, since neither Chi nor Ang, taken singly or in combination, provide any teaching or suggestion of “translation information identifying a translation function”, it follows that the prior art references fail to teach or suggest “[storing] the translation information identifying the translation function” as also recited in claim 1. **Accordingly, Appellant submits that the cited references, taken singly or in combination, fail to teach or suggest “translation information identifying a translation function” or “[storing] the translation information identifying the translation function” as recited in combination with the other features of claim 1.**

In the advisory action, the Examiner contends that Chi teaches a translation function that is translating addresses between global address space and local address spaces of memories being shared in a multi processing clustering system. The Examiner further contends that, in order to carry out this address translation function, in each node, it requires translation information such as Fig. 8 #114 node ID and #110 partition number. However, as discussed above, the information contained in node ID 114 and partition number 110 does not identify a translation function.

Moreover, Fig. 8 and its associated description in Chi is directed to only a single translation function. Chi provides no teaching or suggestion whatsoever that more than a single translation function is implemented in the system disclosed therein. Thus, the translation function disclosed by Chi is the same for all address translations, and thus there is no need to identify a translation function, and thus no need to provide “translation information identifying a translation function” as recited in combination with the other features of claim 1.

Notwithstanding the fact that Chi fails to teach or suggest this limitation (or even provide a rationale for its necessity), Chi further fails to teach or suggest a memory subsystem configured to “**select the translation function in response to the translation information**” as also recited in claim 1. Since Chi discloses only a single translation function and provides no teaching or suggestion of more than one translation function within a single embodiment, Chi neither teaches, nor provides any rationale for providing that the memory subsystem be configured to select a translation function, much less select a translation function in response to translation information. Ang provides no additional teaching or suggestion that would remedy this deficiency, whether taken singly or in combination with Chi. **Accordingly, the cited references, taken singly or in combination, fail to teach or suggest a memory subsystem configured to “select the translation function in response to the translation information.”**

For at least the reasons stated above, Appellant submits that the rejection of claims 1-2, 4-5, and 7-8 is in error, and reversal of the Examiner’s decision is respectfully requested.

B. Claim 3

The Examiner rejected claims 3-4 as being unpatentable over Chi in view of Ang under 35 U.S.C. § 103(a). Appellants respectfully traverse this rejection in light of the following remarks.

Claim 3 depends from claim 1, and is thus believed allowable for at least the same reasons as discussed above. Appellant further submits that the cited references, taken singly or in combination, fail to teach or suggest all of the elements of claim 3 for at least the following additional reasons.

Claim 3 recites:

“The system of claim 1, wherein the additional memory subsystem is configured to perform a **different** translation function on the global address to obtain a local physical address of the coherency unit within the additional memory subsystem” (Emphasis added).

As previously discussed above, Chi is directed to only a single translation function. In contrast, Appellant’s claim 3 is directed to a system having a translation function in a memory subsystem, and a different translation function in an additional memory subsystem. As Chi discloses only the single translation function of Fig. 8, Chi thus does not provide any teaching or suggestion that different translation functions are performed in the various nodes of the system disclosed therein. Furthermore, Ang provides no teaching or suggestion that would remedy this deficiency of Chi. **Accordingly, the cited references, taken singly or in combination, fail to teach or suggest a combination of features including “wherein the additional memory subsystem is configured to perform a different translation function on the global address” as recited in claim 3.**

For at least these reasons, Appellant submits that the rejection of claim 3 is in error, and reversal of the Examiner’s decision is respectfully requested.

C. Claim 9

The Examiner rejected claim 9 as being unpatentable over Chi in view of Ang and Arimili under 35 U.S.C. § 103(a). Appellants respectfully traverse this rejection in light of the following remarks.

Claim 9 depends from independent claim 1 and thus incorporates its limitations. Thus, for at least the reasons stated above regarding claim 1, Appellant submits that the prior art references, taken singly or in combination fail to teach or suggest all of the limitations of claim 9, and respectfully requests reversal of the Examiner's rejection thereof.

D. Claims 10-12, 14, and 16-20

The Examiner rejected claims 10-11, 14, and 16-20 as being unpatentable over Chi in view of Ang under 35 U.S.C. § 103(a). Appellants respectfully traverse this rejection in light of the following remarks.

Independent claim 10 recites:

“A method for use in a system comprising a plurality of nodes, wherein each node includes an active device and a memory subsystem coupled to the active device, the method comprising:

an active device in a node of the plurality of nodes generating a global address and translation information identifying a translation function, wherein the global address identifies a coherency unit;

a memory subsystem included in the node selecting the translation function in response to the translation information and performing the translation function on the global address to generate a physical address of the coherency unit within the memory subsystem;

an additional memory subsystem included in an additional node of the plurality of nodes storing the translation information identifying the translation function used in the node; and

in response to a request for access to the coherency unit, the additional memory subsystem sending the translation information to the node.”

Appellant respectfully submits that the prior art references, taken singly or in combination, fail to teach or suggest “translation information identifying a translation function”, “selecting the translation function in response to the translation information”, “performing [the selected] translation function”, or “storing the translation information identifying the translation function” as recited in claim 10.

In the final office action, the Examiner contends that Chi teaches an active device configured to generate a global address and translation information identifying a translation function and a node configured to perform the translation function in Fig. 8. The Examiner further contends that Chi teaches an additional node configured to store the translation information identifying the translation function in Fig. 8, #114 and #110, and col. 5, lines 27-46. Applicant respectfully disagrees with the Examiner’s characterization. In col. 5, lines 27-46, Chi describes an address translation from the source node to the destination node:

“FIG. 8 shows the address translation from the source node to the destination node. The processor bus address 100 at the source node consists of two parts: an AMT Index 102 and an Offset 104. The AMT Index 102 is used to retrieve an entry from an Address Mapping Table (AMT) 106. Each table entry 108 contains a destination node ID for an access request on the processor bus. The number of entries in the AMT 106 is not necessarily equal to the number of the address partitions. For instance, in one embodiment, the AMT has 64 entries and the number of address partitions is 16. In that case, the AMT Index 102 is 6 bits while only the most significant 4 bits of these 6 bits indicate the partition number 110. In the AMT 106, the four consecutive entries corresponding to a

given partition number are filled with the same destination node. Note that the node ID from the AMT 106 is the physical node ID which is unique in the whole system. The partition number 110 is the logical node ID within each application, such that there may be more than one application running on the system.” (emphasis added).

Chi further states, in column 5, lines 55-65:

“The global address 112 on the interconnect is composed of the node ID 114 and the processor bus address 110 at the source node. This global address is part of the header of the packet sent from the source node to the destination node. At the destination node, the node ID is stripped off and the partition number is replaced with 0, such that the translated address 116 falls in the local space of the destination node. Finally, the private memory size 118 at the destination node (or the starting address of the near global memory of the destination node) is further added to produce the processor bus address 120 to be used at the destination node.” (Emphasis added).

As evidenced by the above two citations, item 110 is taught as either a partition number that is a logical node ID, or a processor bus address. Chi provides no teaching or suggestion that item 110 provides any information that would identify a translation function. As evidenced in the second of the two citations, item 114 is taught by Chi as a node ID. Chi provides no other teaching regarding node ID 114, much less any teaching that node ID 114 identifies a translation function. Furthermore, Chi provides no other teaching or suggestion of “translation information identifying a translation function”, nor does Ang provide any teaching or suggestion that, taken either singly or in combination with Chi that would remedy this deficiency. Finally, since neither Chi nor Ang, taken singly or in combination, provide any teaching or suggestion of “translation information identifying a translation function”, it follows that the prior art references fail to teach or

suggest “storing the translation information identifying the translation function” as also recited in claim 1. **Accordingly, Appellant submits that the cited references, taken singly or in combination, fail to teach or suggest “translation information identifying a translation function” or “storing the translation information identifying the translation function” as recited in combination with the other features of claim 10.**

In the advisory action, the Examiner contends that Chi teaches a translation function that is translating addresses between global address space and local address spaces of memories being shared in a multi processing clustering system. The Examiner further contends that, in order to carry out this address translation function, in each node, it requires translation information such as Fig. 8 #114 node ID and #110 partition number. However, as discussed above, the information contained in node ID 114 and partition number 110 does not identify a translation function.

Moreover, Fig. 8 and its associated description in Chi is directed to only a single translation function. Chi provides no teaching or suggestion whatsoever that more than a single translation function is implemented in the system disclosed therein. Thus, the translation function disclosed by Chi is the same for all address translations, and thus there is no need to identify a translation function, and thus no need to provide “translation information identifying a translation function” as recited in combination with the other features of claim 10.

Notwithstanding the fact that Chi fails to teach or suggest this limitation (or even provide a rationale for its necessity), Chi further fails to teach or suggest a memory subsystem “**selecting the translation function in response to the translation information**” as also recited in claim 10. Since Chi discloses only a single translation function and provides no teaching or suggestion of more than one translation function within a single embodiment, Chi neither teaches, nor provides any rationale for providing that the memory subsystem be configured to select a translation function, much less select a

translation function in response to translation information. Ang provides no additional teaching or suggestion that would remedy this deficiency, whether taken singly or in combination with Chi. **Accordingly, the cited references, taken singly or in combination, fail to teach or suggest a memory subsystem configured to “select the translation function in response to the translation information.”**

For at least the reasons stated above, Appellant submits that the rejection of claims 10-11, 14, and 16-20 is in error, and reversal of the Examiner’s decision is respectfully requested.

E. Claim 12

The Examiner rejected claims 12-13 as being unpatentable over Chi in view of Ang under 35 U.S.C. § 103(a). Appellants respectfully traverse this rejection in light of the following remarks.

Claim 12 depends from claim 10, and is thus believed allowable for at least the same reasons as discussed above. Appellant further submits that the cited references, taken singly or in combination, fail to teach or suggest all of the elements of claim 12 for at least the following additional reasons.

Claim 12 recites:

“The method of claim 10, further comprising the additional memory subsystem performing a **different** translation function on the global address to obtain a local physical address of the coherency unit within the additional memory subsystem.” (emphasis added).

As previously discussed above, Chi is directed to only a single translation function. In contrast, Appellant’s claim 12 is directed to a method wherein a translation function is performed in a memory subsystem, and a different translation function is performed in an

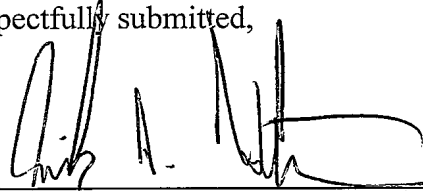
additional memory subsystem. As Chi discloses only the single translation function of Fig. 8, Chi thus does not provide any teaching or suggestion that different translation functions are performed in the various nodes of the system disclosed therein. Furthermore, Ang provides no teaching or suggestion that would remedy this deficiency of Chi. **Accordingly, the cited references, taken singly or in combination, fail to teach or suggest a combination of features including “the additional memory subsystem performing a different translation function on the global address” as recited in claim 12.**

For at least these reasons, Appellant submits that the rejection of claim 12 is in error, and reversal of the Examiner’s decision is respectfully requested.

VII. CONCLUSION

For the foregoing reasons, it is submitted that the Examiner's rejection of claims 1-5, 7-14, and 16-20 was erroneous, and reversal of his decision is respectfully requested.

Respectfully submitted,



Erik A. Heter
Reg. No. 50,652
AGENT FOR APPLICANT(S)

Meyertons, Hood, Kivlin, Kowert, & Goetzel, P.C.
P.O. Box 398
Austin, TX 78767-0398
Phone: (512) 853-8800

Date: 8/3/07

IX. APPENDIX

The claims on appeal are as follows.

1. A system, comprising:

a plurality of nodes, wherein each node includes an active device and a memory subsystem coupled to the active device;

wherein an active device in a node of the plurality of nodes is configured to generate a global address and translation information identifying a translation function, wherein the global address identifies a coherency unit;

wherein a memory subsystem included in the node is configured to select the translation function in response to the translation information and to perform the translation function on the global address to generate a physical address of the coherency unit within the memory subsystem;

wherein an additional memory subsystem included in an additional node of the plurality of nodes is configured to store the translation information identifying the translation function used in the node, wherein in response to a request for access to the coherency unit, the additional memory subsystem is configured to send the translation information to the node.

2. The system of claim 1, wherein the plurality of nodes are coupled by an inter-node network, and wherein each of the plurality of nodes includes an interface to the inter-node network;

wherein an additional interface included in the additional node is configured to receive the translation information for the coherency unit from the additional memory subsystem;

wherein the additional interface is configured to provide the translation information and the global address to an interface included in the node via the inter-node network; and

wherein the interface included in the node is configured to provide the translation information and the global address received via the inter-node network to the memory subsystem.

3. The system of claim 1, wherein the additional memory subsystem is configured to perform a different translation function on the global address to obtain a local physical address of the coherency unit within the additional memory subsystem.
4. The system of claim 1, wherein the additional memory subsystem is configured to store translation information associated with the coherency unit for several other nodes included in the plurality of nodes, wherein different translation information is associated with each of the several other nodes.
5. The system of claim 1, wherein each active device included in the plurality of nodes is configured to use at least a portion of the global address of the coherency unit to determine whether a copy of the coherency unit is locally cached by that active device.
7. The system of claim 1, wherein no memory subsystem included in an other node of the plurality of nodes maps the global address, wherein an other active device included in the other node is configured to request access to the coherency unit by sending a packet including the global address and translation information

associated with the coherency unit in the other node on a network included in the other node, wherein the translation information associated with the coherency unit in the other node indicates that no memory subsystem included in the other node maps the coherency unit.

8. The system of claim 7, wherein an other interface included in the other node and coupled to the network is configured to forward the global address to an additional interface included in the additional node in response to receiving the packet.
9. The system of claim 1, wherein a memory controller included in the memory subsystem is integrated in a same integrated circuit as the processing subsystem.
10. A method for use in a system comprising a plurality of nodes, wherein each node includes an active device and a memory subsystem coupled to the active device, the method comprising:

an active device in a node of the plurality of nodes generating a global address and translation information identifying a translation function, wherein the global address identifies a coherency unit;

a memory subsystem included in the node selecting the translation function in response to the translation information and performing the translation function on the global address to generate a physical address of the coherency unit within the memory subsystem;

an additional memory subsystem included in an additional node of the plurality of nodes storing the translation information identifying the translation function used in the node; and

in response to a request for access to the coherency unit, the additional memory subsystem sending the translation information to the node.

11. The method of claim 10, further comprising:

an additional interface included in the additional node receiving the translation information for the coherency unit in the node from the additional memory subsystem;

the additional interface providing the translation information and the global address to an interface included in the node via an inter-node network; and

the interface included in the node providing the translation information and the global address received via the inter-node network to the memory subsystem.

12. The method of claim 10, further comprising the additional memory subsystem performing a different translation function on the global address to obtain a local physical address of the coherency unit within the additional memory subsystem.

13. The method of claim 10, further comprising the additional memory subsystem storing translation information associated with the coherency unit for several other nodes included in the plurality of nodes, wherein different translation information is associated with each of the several other nodes.

14. The method of claim 10, further comprising each active device included in the plurality of nodes using at least a portion of the global address of the coherency unit to determine whether a copy of the coherency unit is locally cached by that active device.

16. The method of claim 10, further comprising an other active device included in an other node of the plurality of nodes requesting access to the coherency unit by sending a packet including the global address and translation information associated with the coherency unit in the other node on a network included in the other node, wherein no memory subsystem included in the other node maps the global address, and wherein the translation information associated with the coherency unit in the other node indicates that no memory subsystem included in the other node maps the coherency unit.
17. The method of claim 16, further comprising an other interface included in the other node and coupled to the network is configured to forward the global address to an additional interface included in the additional node in response to receiving the packet.
18. The method of claim 10, wherein a memory controller included in the memory subsystem is integrated in a same integrated circuit as the processing subsystem.
19. The method of claim 10, further comprising an operating system executing on the active device creating a translation lookaside buffer entry corresponding to the virtual address, wherein the translation lookaside buffer entry includes the global address and the information identifying the translation function, wherein the operating system selects the translation function in order to map the global address to the local physical address within the memory.
20. The method of claim 10, further comprising an operating system executing on the active device in the node creating a translation lookaside buffer entry corresponding to a virtual address in response to deciding to replicate the coherency unit to the node from the additional node, wherein the translation

lookaside buffer entry corresponding to the virtual address specifies the global address and the information identifying the translation function.

X. EVIDENCE APPENDIX

No evidence submitted under 37 C.F.R. §§ 1.130, 1.131, or 1.132 or otherwise entered by the Examiner is relied upon in this appeal.

XI. RELATED PROCEEDINGS APPENDIX

There are no related proceedings.